arXiv:0803.1628v1 [stat.ML] 11 Mar 2008

Component models for large networks

Janne Sinkkonen JANNE.SINKKONEN@TKK.FI
Xtract Ltd. and Helsinki University of Technology

Janne Aukia JANNE.AUKIA@XTRACT.COM
Xtract Ltd.
Hitsaajankatu 22, 00810 Helsinki, Finland

Samuel Kaski SAMUEL.KASKI@TKK.FI
Department of Information and Computer Science
Helsinki University of Technology
P.O. Box 5400, FI-02015 TKK, Finland

Abstract

Being among the easiest ways to find meaningful structure from discrete data, Latent Dirichlet Allocation (LDA) and related component models have been applied widely. They are simple, computationally fast and scalable, interpretable, and admit nonparametric priors. In the currently popular field of network modeling, relatively little work has taken uncertainty of data seriously in the Bayesian sense, and component models have been introduced to the field only recently, by treating each node as a bag of out-going links. We introduce an alternative, interaction component model for communities (ICMc), where the whole network is a bag of links, stemming from different components. The former finds both disassortative and assortative structure, while the alternative assumes assortativity and finds community-like structures like the earlier methods motivated by physics. With Dirichlet Process priors and an efficient implementation the models are highly scalable, as demonstrated with a social network from the Last.fm web site, with 670,000 nodes and 1.89 million links.

Keywords: Latent-Component Mixture Model, Social Network, Probabilistic Community Finding, Nonparametric Bayesian

1. Introduction

Data collections representable as networks, or sets of binary relations between vertices, appear now frequently in many fields, including social networks and interaction networks in biology (Fig. 1). Consequently, inferring properties of the network vertices[1] from the edges has become a common data mining problem. Most of the work has been about dividing the vertices into relatively well-connected subsets, or *communities* (Fortunato and Castellano, 2007). Most papers on communities have been inspired by graph theory and physics, as is a large field of fundamental network-related work not directly relevant here. Especially optimizing a measure of good division called *modularity* (Newman, 2006) has gained success but is not without its problems (Fortunato and Barthelemy, 2007; Kumpula et al., 2007).

1. We will use the terms vertex and node interchangeably, and likewise for edges and links.

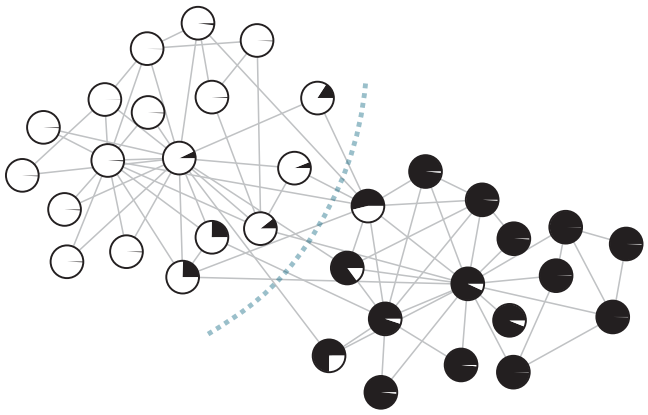

Figure 1: One of the models (ICMc) on a classic small network, the Karate club. Zachary (1977) observed the social interactions (edges) between 34 members of a karate club over two years. During this period, there was disagreement among the club members which eventually led to the splitting of the club (dashed line). Black and white depict the degree of component memberships obtained by the model, without knowing the correct split.

A feature and potential problem of modularity is that it takes the observed edges granted, while network data are typically not a complete description of reality but comes with errors, omissions and uncertainties. Some links may be spurious, for instance due to measurement noise in biological networks, and some potential links may be missing, for instance friendship links of newcomers in social networks. Probabilistic generative models are a tool for modeling and inference under such uncertainty. They treat the links as random events, and give an explicit structure for the observed data and its uncertainty. Compared to non-stochastic methods, they are therefore likely to perform well as long as their assumptions are valid: They may reveal properties of networks that are difficult to observe with non-statistical techniques from the noisy and incomplete data, and they also offer a groundwork for new conceptual developments. For example, it may be argued that network communities should be defined in terms of stochastic models that do not take links at face value but instead give them an underlying stochastic structure that should be realistic given an application. On the down side, probabilistic methods are not always scalable, and they may be difficult to understand, apply and trust by people from other fields, especially if the estimation process is complex.

Probabilistic models of network connectivity have been introduced recently. Mixtures of latent components (Newman and Leicht, 2006), analogous to finite mixture models for vectorial data, are attractive because of ease of interpretation, but the extensive numbers of parameters encumber straightforward fitting attempts. A very promising development called stochastic block models (Airodi et al., 2008—but also Daudin et al., 2007; Hofman and Wiggins, 2007) groups the nodes into blocks and explains the links in terms of homogeneous connections between pairs of groups. Finally, links can be explained by the proximity of

nodes in a latent space created by a logistic link (Handcock et al., 2007). These models have been successively applied to various networks from sociology and biology, up to the size of thousands or tens of thousands of nodes. With heuristic improvements, stochastic block models are expected to scale up to over one million nodes (E. Airoldi, p.c.), but in general the computational bottleneck is scalability.

The models discussed in this paper are generative probabilistic models that decompose the links into components, but their structure makes them scalable to networks with at least $10^3 \ldots 10^6$ nodes, and up to thousands of latent components—as long as the networks are sparse enough. The Simple Social Network LDA (SSN-LDA) model presented by Zhang et al. (2007) is identical to the Latent Dirichlet Allocation (LDA; Buntine, 2002; Blei et al., 2003) model, originally applied to text collections. It is also a conceptual although not a geneologic successor of the mixture model by Newman and Leicht (2006). The SSN-LDA model assumes that each node is a bag of outgoing links, and models each outgoing set of links as a mixture over latent components. The components are the same for each node, but their proportions differ.

As an alternative we introduce a component model for relational data, where each link is directly assumed to come from a latent component, and the whole network is a bag of links (Sinkkonen et al., 2007). This model is particularly well suited for modeling of community-type structure in networks. For conciseness, we call it ICMc (interaction component model for communities), the latter 'c' reminding of the fact that it is easy to generate new models from the family of ICMc and SSN-LDA, with slightly different generative assumptions and requirements for data.

Both ICMc and SSN-LDA represent a set of links as a probabilistic mixture over latent components. Depending on the prior, the models can find either a given number of latent components, or nonparametrically adjust the number of components to the data, guided by a diversity parameter. Moreover, depending on parameters, they are capable of finding either subnetworks or more graded, latent-space-like structures.

Both models can be easily and efficiently fitted to data by collapsed Gibbs sampling (Neal, 2000), an MCMC technique for sampling from the posterior where parameters have been integrated out and latent variables are sampled. In the component models the latent variables give the assigments of the links to the components. Critical for successful scaling to large networks is sparseness of representations; here the component assignments of the links, the variables that are sampled in the collapsed Gibbs, can be efficiently represented as sparse arrays, trees, and hash maps.

We compare the two models on two citation networks with a few thousand nodes, CiteSeer and Cora (Sen and Getoor, 2007), and demostrate their properties on smaller networks. As a demonstration of a larger-scale problem, musical tastes of people are derived from the friendship network of the online music service Last.fm (`www.last.fm`), with over 650,000 vertices vertices and almost two million edges.

2. Two scalable network models

SSN-LDA models directed links. A unique mixing pattern over latent *link target profiles* is associated to each node. (Technical details are presented later, e.g., in Fig. 4, right). The latent profiles correspond to *topics* of text document models, the original application

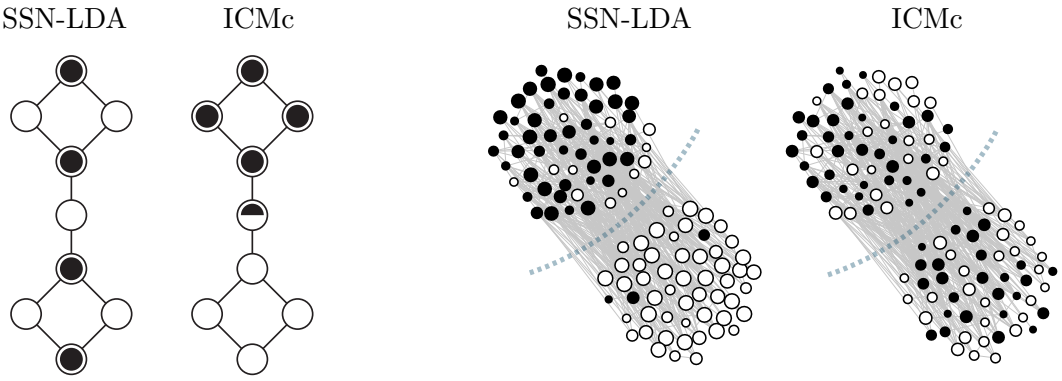

Figure 2: Network structure and component models. *Left:* A toy network split into two components (black and white). SSN-LDA produces both assortative and disassortative components, but here favors the disassortative interpretation, grouping nodes by common external connectivity. The ICMc solution is assortative and more community-like. *Right:* The Adj-Noun network of adjective and noun co-occurrences is bipartite, with a negative modularity of -0.241. SSN-LDA, not limited to assortative solutions, finds the underlying structure (node size represents certainty). Modularity of the SSN-LDA solution (black vs. white) is -0.262 and that of ICMc 0.188. See Table 1 for details of Adj-Noun.

of LDA. If the node memberships in latent profiles are sharp enough, that is, if the nodes are mainly associated to one profile only, the profiles can be interpreted as subgraphs. The grouping criterion is a probabilistic version of the *structural equivalence* principle of sociology (Michaelson and Contractor, 1992): Two nodes belong to the same group if their role in the network topology is similar, that is, they link to the same (other) nodes.

In ICMc, a unique mixture over latent components is associated with each node, and linking is unstructured inside a component. Instead of structural equivalence, the criterion for subgroups is homogeneous, symmetric internal connectivity. Link directions are therefore not modelled. A related social concept is subgroup cohesion (Wasserman and Faust, 1994), where latent similarity results in connections inside the group, instead of linking into some common third party. As a result, the network looks *homophilic* (Lazarsfeld and Merton, 1954); the connected nodes tend to be relatively similar by their non-network properties.

For technical reasons, the parameterization of linking within a component in ICMc is in terms of linking probabilities over the components; memberships of nodes in components can be obtained from these parameters by the Bayes rule. Equivalently, the model can be described as modeling the whole graph as a bag of links. Each link comes from a component specified by a latent variable z (Fig. 4, left). Each component chooses the endpoints of a link from a component-specific (multinomial) distribution over the nodes, parameterized by m_z.

A further helpful distinction is that of *assortative* and *disassortative* network properties. A network is assortative with respect to a property if the property tends to co-occur in connected nodes more often than expected by change (Newman, 2003). The opposite,

negative correlation in adjacent nodes, is called disassortativity. SSN-LDA can in principle find either kinds of structures, while ICMc tends to find only assortative structure.[2] Modularity, a quality measure for community detection (Newman, 2006), can at least to a degree be used as a measure of assortativity; if it is negative for a partitioning of the network, the partitioning is disassortative (Fortunato and Castellano, 2007). Unfortunately, comparing modularities of partitionings over different networks is in general not justified (Fortunato and Castellano, 2007), and hence we cannot use it to compare the modeling problems.

One would expect ICMc to find communities from social and other networks better than the less specialized SSN-LDA, as long as linking results from homophily and the communities can be assumed assortative. The reason is that a model having less degrees of freedom in its parameterization will be able to more accurately estimate the parameters from the relatively small observed data sets. On the other hand, ICMc should be unable to find disassortative structures. This seems to indeed hold in some extreme cases (Fig. 2), but in practice differences are often graded and harder to attribute to properties of the network (Fig. 3).

The behavior of both SSN-LDA and ICMc are determined by their hyperparameters. Both models can be made to prefer either latent components of equal size, or to allow heavy size variation. Even more importantly, either graded or non-overlapping components can be preferred. In graded components the community membership probabilities are akin to coordinates in a latent space, while non-overlapping components divide the nodes sharply into clusters.

Both models can accommodate integer link weights in the sense of generating multiple links between two nodes. On the other hand, the models work particularly efficiently for sparse binary links: If data is sparse, link probabilties are small overall, multiple links even more improbable, and the model effectively generates binary data.

The SSN-LDA model is originally based on a finite mixture (Zhang et al., 2007), but it is easily extended for a Dirichlet process prior (DP prior; Blackwell and MacQueen, 1973; Neal, 2000), while ICMc is originally with a DP prior (Sinkkonen et al., 2007), but here applied also in its finite form.

The models are demonstrated on three small networks in Figures 1 and 3. The first is the Karate network originating from a study by Zachary (1977). In the study Zachary observed the social interactions between 34 members of a karate club over two years. During this period, there was disagreement among the club members which led to the splitting of the club. Figure 1 demonstrates that ICMc finds the splitting.

The second demonstration network is the Football network (Girvan and Newman, 2002), which depicts American football games between Division IA colleges during the fall season 2000. The nodes of the network represent football teams and edges the games between the teams. There is a known community structure for the network in the form of conferences. In general, games between teams that belong to the same conference are more frequent than games between teams that belong to different conferences, but sometimes teams prefer to play mostly against teams in other conferences. Both models find the structure as seen in Figure 3, SSN-LDA slightly more accurately. ICMc is somewhat more accurate on another network derived from political blogs.

2. The discussion of the distinction by Newman and Leicht (2006) is indeed applicable to SSN-LDA, for SSN-LDA can be seen as a Bayesian extension of the earlier model.

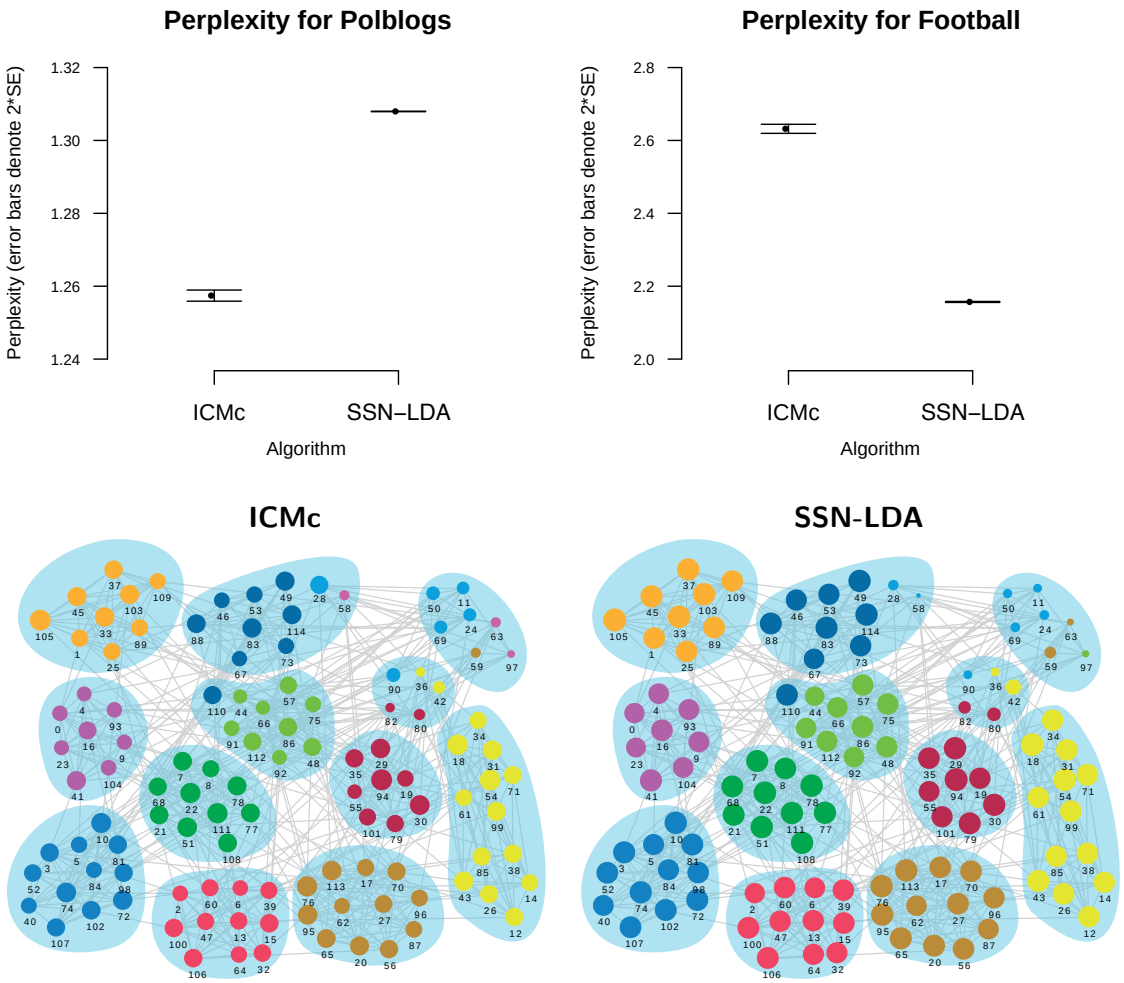

Figure 3: Relative performance of ICMc and SSN-LDA varies for different networks. *Above:* ICMc performs better for a network of political blogs, as measured by the perplexity of the components, while SSN-LDA is better for a network of US football games. Reasons for the differences are unclear, although they might be related to the assortativity of the networks with respect to the ground clusters (political orientation and football conferences). *Below:* Despite perplexity differences, the solutions are qualitatively very similar. For the Football network, main differences are in the certainty of cluster assignments. Shaded areas show the borders of the conferences, while community assignments by the model and their certainty are depicted by node color and size, respectively. See Table 1 for details of the networks and model parameters.

Table 1: Network characteristics and modeling parameters for the small and medium-size networks. In the table, I is the number of nodes in the network, L is the number of edges, Q the ground cluster modularity, and α_{Dir} and β are the hyperparameters of the models.

				ICMc		SSN-LDA	
Network	I	L	Q	α_{Dir}	β	α_{Dir}	β
Adj-Noun	112	423	-0.241	0.5	0.2	0.5	0.2
Football	115	613	0.554	0.083	0.03	0.083	0.7
Polblogs	1 222	16 714	0.410	0.5	0.003	0.5	0.4
Citeseer	2 120	3 678	0.517	0.166	0.04	0.166	0.006
Cora	2 485	5 067	0.630	0.143	0.02	0.143	0.025

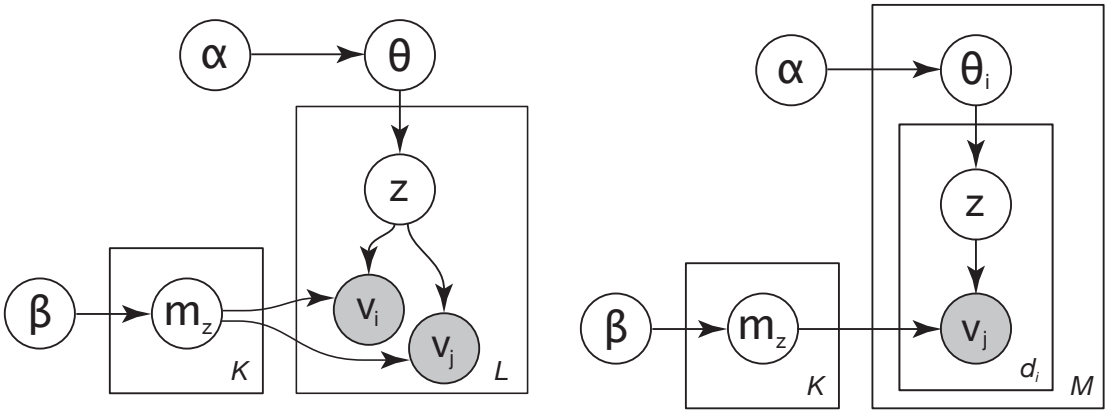

Figure 4: The ICMc model (*left*) and SSN-LDA (*right*). SSN-LDA is effectively the Latent Dirichlet Allocation (LDA; Buntine, 2002; Blei et al., 2003) applied to network data, with nodes playing both the role of 'words', at the receiving end of links, and 'documents' at the sending end. ICMc has no hierarchy level for nodes. Instead, it generates *two* nodes for each link; the links are undirected. See Section 2 for the notation and further discussion.

2.1 Interaction Component Model for Communities (ICMc)

The generative process out of which the network is supposed to arise is the following (see Fig. 4 for a diagram); it is parameterized by the hyperparameters (α, β).

(1.1) Generate a multinomial distribution θ over *latent components* z. For K components, the multinomial is generated from a K-dimensional Dirichlet distribution with all parameters set to α_{Dir}, $\mathrm{Dir}^K_{\mathrm{sym}}(\alpha_{Dir})$, and for an infinite number of components from the Dirichlet process $\mathrm{DP}(\alpha_{DP})$.

(1.2) To each z, associate a multinomial distribution over the M vertices i by sampling the multinomial parameters m_z from the Dirichlet distribution $\mathrm{Dir}^M_{\mathrm{sym}}(\beta)$. (To clarify, we have $\sum_i m_{zi} = 1$ for each z, and $\sum_z \theta_z = 1$.);

(2) Then repeat for each link $l = 1 \ldots L$:

(2.1) Draw a latent component z from the multinomial θ.

(2.2) Choose two nodes, i and j, independently of each other, with probabilities m_z; set up a nondirectional link between i and j.

Within components, edges are generated independently of each other; the non-random structure of the network emerges from the tendency of components to prefer certain vertices (that is, m). In contrast to many other network models, the latent variables operate on the edge level, not on the vertex level. There is no explicit hierarchy level for vertices, but because vertices typically have several edges, they are implicitly treated as mixtures over the latent components. Finally, the model is parameterized to generate self-links and multi-edges because this choice allows sparse implementations which would not be directly possible with a potential alternative model that would generate binary links from the Bernoulli distribution.

Although in the case of a Dirichlet process prior the number of potentially generated components is infinite, the prior gives an uneven distribution over the components. Therefore, with a suitably small value of α_{DP}, we observe much fewer components than the number of links is, and the model is useful. On the other hand, β describes the unevenness of the degree distribution of the nodes *within components*: a high β tends to give components spanning over all nodes, while a small β prefers mutually exclusive, community-like components.

We have estimated the model with Gibbs sampling (Geman and Geman, 1984), a variant of MCMC methods that produce samples from the posterior distribution of the model parameters and the latent component memberships. As a side note, maximum likelihood or MAP estimation of the model is not sensible since the number of parameters and latent variables is large compared with the available data, It is easy to derive an EM algorithm for the finite-mixture ICMc, but it gets stuck into suboptimal local posterior maximums at the borders of the parameter space.

We use Gibbs sampling with some of the model parameters integrated out, called Rao-Blackwellized, or collapsed (Neal, 2000). (For the joint distribution of the model and the derivation of the estimation algorithm see the Appendix). In the collapsed Gibbs estimation algorithm the unknown model parameters m_{zi} and θ_z are marginalized away and only the latent classes of the edges, z_l, are sampled, one edge at a time. In general, we denote edge counts per component by n_z, component-wise vertex degrees by k_{zi}, and the endpoints of the left-out edge by (i, j). Then delete one edge, resulting in counts (n', k', N') that are equal to (n, k, N) but without the one edge. The component probabilities of the left-out edge are

$$p(z|i,j) \propto \frac{k'_{zi} + \beta}{2n'_z + 1 + M\beta} \times \frac{k'_{zj} + \beta}{2n'_z + M\beta} \times \frac{n'_z + \alpha_{Dir}}{N' + K\alpha_{Dir}} \tag{1}$$

Algorithm 1 ICMc-Gibbs. A simple implementation for the ICMc algoritm with a Dirichlet prior.

SIMPLE-GIBBS-SAMPLING(α_{Dir}, β, L)

$t_{nodes} \leftarrow$ node count
for $c \leftarrow 1$ **to** *components* **do** ▷ initialize data structures
  $A[c] \leftarrow 0$
  for $n \leftarrow 1$ **to** t_{nodes} **do** $K[n, c] \leftarrow 0$
for $i \leftarrow 1$ **to** *iterations* **do** ▷ main iteration loop
  foreach l **in** L **do**
    $v_i, v_j \leftarrow$ first and second node of l
    if $i \neq 1$ **do**
      $z_{old} \leftarrow Z[l]$
      decrement $K[v_i, z_{old}], K[v_j, z_{old}], A[z_{old}]$
    for $c \leftarrow 1$ **to** *components* **do**
      $P[c] \leftarrow$ CALC-PROBABILITY($A[c], K[v_i, c], K[v_j, c], t_{nodes}, \alpha_{Dir}, \beta$)
    $z_{new} \leftarrow$ sample index from P
    $Z[l] \leftarrow z_{new}$
    increment $K[v_i, z_{new}], K[v_j, z_{new}], A[z_{new}]$
return K, Z

CALC-PROBABILITY($n_c, k_a, k_b, t_{nodes}, \alpha_{Dir}, \beta$)

return $\dfrac{(k_a + \beta)(k_b + \beta)(n_c + \alpha_{Dir})}{(2n_c + 1 + \beta t_{nodes})(2n_c + \beta t_{nodes})}$

for the Dirichlet prior, and

$$p(z|i,j) \propto \frac{k'_{zi} + \beta}{2n'_z + 1 + M\beta} \times \frac{k'_{zj} + \beta}{2n'_z + M\beta} \times \frac{\mathrm{C}(n'_z, \alpha_{DP})}{N' + \alpha_{DP}} \tag{2}$$

for the Dirichlet process prior. The 'chooser' function $\mathrm{C}(n_z, \alpha_{DP}) \equiv n_z$ if $n_z \neq 0$ and $\mathrm{C}(0, \alpha_{DP}) = \alpha_{DP}$. The case with α_{DP}, as opposed to n_z, corresponds to a new component with no other links so far.

This sampling step is simply repeated iteratively for all links, until convergence to the posterior distribution, or until the results are satisfactory by some other measure. A particularly elegant, although not necessarily the most efficient initialization of the sampler starts from empty urns, with $k_{zi} = n_z = N = 0$, then runs through the edges once in a random order and populates the urns according to (1) or (2) while counting only the edges seen so far.

The goal of model fitting is usually to infer community memberships of the nodes. From the Bayes rule we obtain

$$p(z|i) = \frac{\theta_z m_{zi}}{\sum_{z'} \theta_{z'} m_{z'i}} \,. \tag{3}$$

A sample of the marginalized parameters θ and m can be reconstructed from each realization of the counts (k, n) by sampling from the conditional Dirichlet distributions given the priors and the counts:

$$\theta \sim \underset{z}{\mathrm{Dir}}(n_z + \alpha_{Dir}) \text{ or } \theta \sim \underset{z}{\mathrm{Dir}}(\{n_z, \alpha_{DP}\}_z), \text{ and } m_z \sim \underset{i}{\mathrm{Dir}}(k_{zi} + \beta) . \tag{4}$$

Note that $n_z = \sum_i k_{zi}/2$. In the case of the Dirichlet process prior, the parameter θ has probabilities of the components with at least one assigned link, and then the probability of all empty components summed up into the last bin. These correspond to the Dirichlet parameters n and α_{DP}, respectively.

Even if one wants to reconstruct θ and m, collapsed Gibbs is likely to be faster than full Gibbs. The reasons are twofold: Firstly, Gibbs converges faster when the parameter updates are not in the main loop. Secondly, one usually uses decimation in sampling from the converged chain, and the (θ, m) need to be constructed only for the decimated samples.

It is often sufficient to estimate the community memberships from the expected values of the marginalized parameters,

$$\hat{\theta} = \frac{n_z + \alpha_{Dir}}{\sum_{z'} n_{z'} + K\alpha_{Dir}} \text{ or } \hat{\theta} = \frac{n_z}{\sum_{z'} n_{z'} + \alpha_{DP}} , \tag{5}$$

and

$$\hat{m}_z = \frac{k_{zi} + \beta}{\sum_{i'} k_{zi'} + M\beta} . \tag{6}$$

Substituting the expectations into (3), we find that for small α and β,

$$p(z|i) \approx \frac{k_{zi}}{\sum_{z'} k_{z'i}} \tag{7}$$

is a good approximation.

Prediction for new data is straightforward; the component memberships of the links associated to a new node can be sampled from (1), given old links. If the new links are not conditionally independent given the old data, one can run a short Gibbs iteration on the new links.

2.2 SSN-LDA

SSN-LDA (Zhang et al., 2007) also has two hyperparameters, denoted by α and β, but they are in a slightly different role than in ICMc (see Fig. 4). The generative process is as follows:

(1.1) Generate M multinomial distributions θ_i, $i = 1, \ldots, M$, over *latent components* z, $z = 1, \ldots, K$, either from a K-dimensional Dirichlet distribution $\mathrm{Dir}^K_{\mathrm{sym}}(\alpha_{Dir})$, or from the Dirichlet process $\mathrm{DP}(\alpha_{DP})$.

(1.2) Assign a multinomial distribution m_z over the vertices i to each component z by sampling from the Dirichlet distribution $\mathrm{Dir}^M_{\mathrm{sym}}(\beta)$.

(2) Then repeat for each link $l = 1, \ldots, L$, with sending nodes $i = 1, \ldots, M$:

(2.1) Draw a latent component z from the multinomial θ_i.

(2.2) Choose the link endpoint j with probabilities m_z; set up a directional link between i and j.

We have presented the generative process of links in a flat form to make comparison to ICMc easier; in step 2, the loop over nodes is avoided by referring to the node indices i associated to links.

In contrast to ICMc, SSN-LDA has the node as an explicit hierarchy level—in the generative model, there are the parameters θ for each node separately, and α is the common hyperparameter of these node-wise distributions. As in ICMc, the hyperparameters m are associated to latent components over nodes, and β determines their prior. But now m determines only the probabilities of the *receiving nodes*. Sending probabilities, associated to starting points i of the links, are modeled by θ.

Collapsed Gibbs sampling operates on two sets of counts: n_{iz} that counts the sender–component combinations (i, z) for links, originating from step (2.1) of the generative process, and k_{zj} counting the receiver–component combinations (j, z) from step (2.2) of the process. Following Griffiths and Steyvers (2004), and Zhang et al. (2007), the conditional probabilities for sampling a left-out link in a collapsed Gibbs iteration, given hyperparameters and all other links, is

$$p(z|i,j) \propto \frac{k'_{zj} + \beta}{k'_{z\cdot} + M\beta} \times \frac{n'_{iz} + \alpha_{Dir}}{n'_{i\cdot} + K\alpha_{Dir}} \,, \tag{8}$$

where sums over counts k' and n' have been denoted by the dot notation. We have omitted the derivation of the Dirichlet process variant, because it is very similar to the derivation of the DP ICMc (see the Appendix), leading to:

$$p(z|i,j) \propto \frac{k'_{zj} + \beta}{k'_{z\cdot} + M\beta} \times \frac{\mathrm{C}(n'_{iz}, \alpha_{DP})}{n'_{i\cdot}} \,. \tag{9}$$

Again, parameter reconstruction for θ and m can be done either by sampling from the corresponding Dirichlet distributions, or by computing the conditional MAP estimates, either roughly or exactly including priors. As with ICMc, we have used the rough alternative suitable for small values of α and β:

$$p(z|i) \approx \frac{n_{iz}}{n_{i\cdot}} \,. \tag{10}$$

This is for the community memberships of nodes as senders of links. Because in SSN-LDA links are directed, it is possible to define the memberships also in terms of received links,

$$p(z|j) \approx \frac{k_{zj}}{k_{\cdot j}} \,. \tag{11}$$

2.3 Efficient implementation of the collapsed Gibbs samplers

Large real-life networks are sparse almost by definition, and for efficiency it is important to preserve the sparseness in model structures. ICMc and SSN-LDA facilitate sparse structures, since likelihoods decompose into sums over existing links, and terms related to non-links do not appear.

Collapsed Gibbs sampling of ICMc and SSN-LDA needs tables for n and k which together, as a first approximation, are of the size complexity $O(MK)$. In addition one needs to keep track of the component identities of the links, an array of size $O(L)$. But in both models the degree of a node poses an upper limit for its component heterogeneity, so that only a few of the counts k, or k and n in LDA, are simultaneously non-zero, allowing sparse representation of the count tables. Therefore with hash tables memory consumption can be reduced to $O(M\bar{d} + L + K)$ where $\bar{d}$ is the average degree of a node. Because $\bar{d} = L/M$, memory consumption scales as $O(L + K)$.

Marginal sums of the count tables, notably n in ICMc, can be represented in a sparse form and updated efficiently during sampling with the aid of a self-balancing binary tree. The idea of using a tree in sampling of discrete distributions was originally proposed by Wong and Easton (1980), and another method for using binary trees in simulations is provided by Blue et al. (1995). In our implementation the Arne-Andersson tree (AA tree) is used (see, e.g., Weiss, 1998), but other self-balancing binary trees would be equivalent in performance. A partial sum tree is formed, where in each node, the total probability of the node is stored, together with the sums of probabilities of both the left and right children of the node. When the probability of a node is changed, the modifications are propagated up to the parents of the node. Sampling proceeds recursively down the tree as a sequence of weighted Bernoulli samples.

These sparse representations, and the binary tree for the marginal sums, make it possible to run models with at least tens of thousands of components in an ordinary PC or server. These structures also fit well with the dynamic component numbers due to the Dirichlet process prior. With the data structures described above, running time per one Gibbs iteration over all the nodes becomes $O(L\,\bar{d}\,\log K)$. That is, the time needed for an iteration scales linearly in the number of edges and logarithmically in the number of components.

It is hard to give any general rule on the number of Gibbs iterations needed for convergence. Because the variables in the collapsed Gibbs algorithm correspond to links, the dependency graph of the variables is like the original network, but with the nodes being in the role of links, and vice versa. The path lenghts of the dependency network are therefore proportional to the path lengths of the original network. Let us assume that the average path length scales as $l \propto \log M$, as is the case with many small-world networks (Albert and Barabasi, 2002). In Gibbs, information diffusion over the network can be expected to take l^2 iterations, analogously to ordinary diffusion. This leads to the conjecture that the number of Gibbs iterations should be proportional to $\log^2 M$.

3. Tests

We compared SSN-LDA and ICMc on two medium-scale social network datasets, Cora and Citeseer (Sen and Getoor, 2007), in the task of finding a predefined set of known clusters. Performance on large networks of $10^5 \ldots 10^6$ nodes is then demonstrated for one of the models (ICMc) with two friendship networks from the music site Last.fm.

3.1 ICMc vs. SSN-LDA

The Cora and the CiteSeer datasets consist of content descriptions of scientific publications and citations between them. The Cora dataset has 2,708 papers in seven predefined classes,

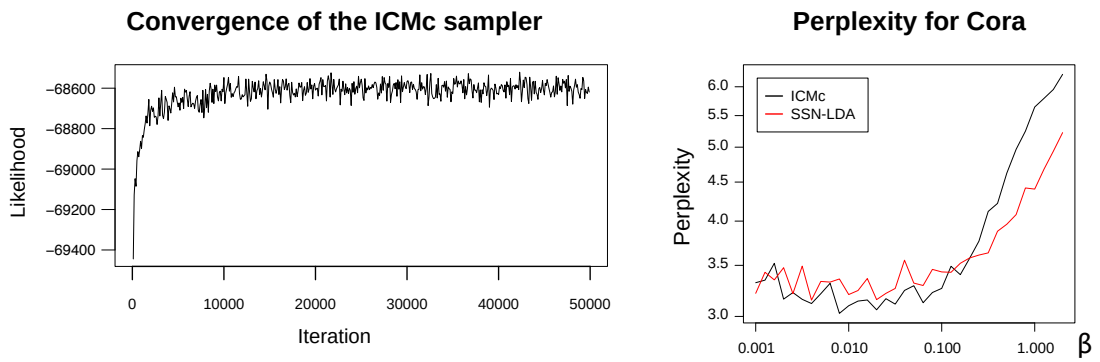

Figure 5: Gibbs samplers on the Cora citation set: convergence and sensitivity to the hyperparameter β. *Left:* Leave-one-out logarithmic posterior probability of the data for a single ICMc chain. This can be recorded easily during sampling as log probabilities of the drawn link assignments. We ran 50,000 iterations over the data, but about 15,000 would have been enough for convergence, and about 3,000 for getting useful results. SSN-LDA convergence was very similar. *Right:* Perplexity for the Cora dataset with a range of hyperparameter values β. Each reported value is an average of four chains. Both models are quite robust with respect to β.

while the CiteSeer dataset contains 3,312 publications in six classes. We used only the citation information, and the predefined classes as a ground truth for clustering. Nodes (publications) not belonging to the main components of the network were removed, and directional links were symmetrized. The resulting network for Cora has 2,120 nodes and for Citeseer 2,485 nodes (Table 1).

Following Zhang et al. (2007) and our own experiences (Aukia, 2007), we fixed $\alpha_{Dir} = 1/K$ for both models and datasets. Values for the parameter β were chosen with pretests (Table 1 and Fig. 5). In general, the models with a Dirichlet prior and a small number of components are quite insensitive to values of α_{Dir} and β within the range $0.001 \ldots 0.1$.

The Gibbs sampler was initialized as suggested in Section 2.3 and run for 50,000 iterations (see Fig. 5). We then took 100 samples at intervals of 100. Each sample consists of the latent cluster memberships z for all links. Node memberships were constructed by (7) and (10) for each sample separately, and these were summed up to get confusion matrices.

Over the computed 50 chains, there is a good average correspondence between the found clusters and the original manual clustering of the data sets (Fig. 6). In terms of perplexity ICMc is able to recover the orignal clusters better than SSN-LDA, although the average confusion matrices are relatively similar. Results vary from chain to chain more than with small networks, indicating multiple local minima for the Gibbs sampler to get trapped into. (See Section 4 below for discussion on this behaviour.)

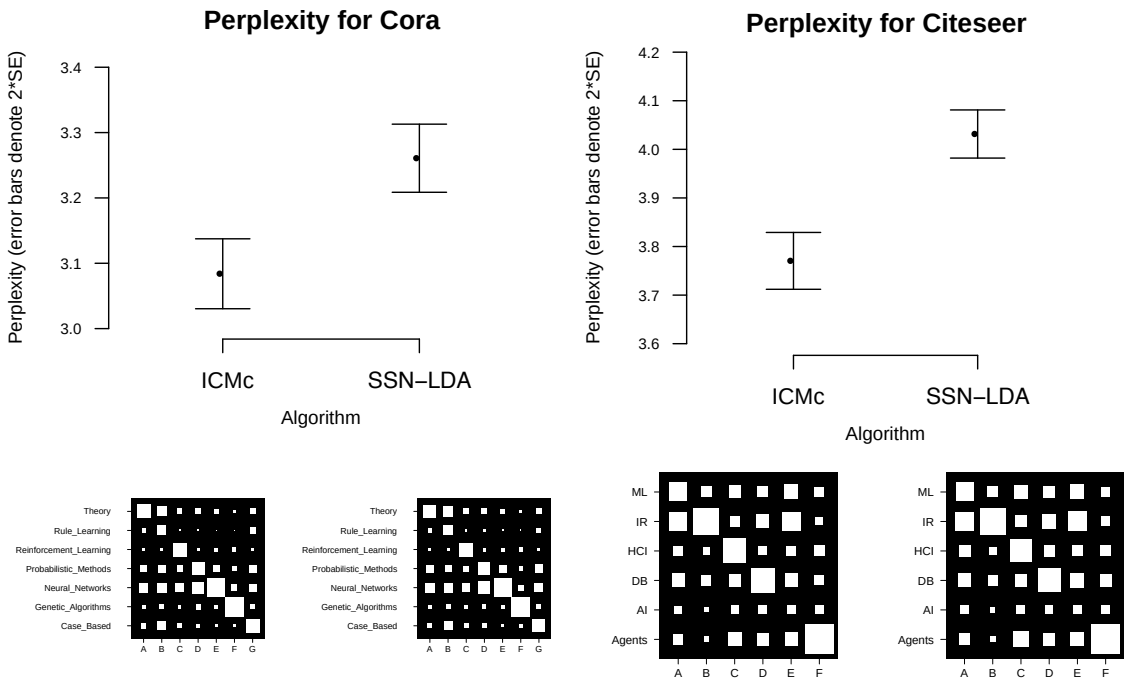

Figure 6: The models ICMc and SSN-LDA on two citation networks, Cora and CiteSeer. *Above:* Performance in finding true clusters, as measured by the perplexity of predicting ground-truth groups with the clusters. The average and 95% confidence intervals for the mean are over 50 chains. *Below:* Average confusion matrices between the found clusters (columns) and the true clusters (rows).

Table 2: Last.fm networks and modeling parameters: I is the number of nodes in the network, L is the number of edges, and α_{DP} and β are the hyperparameters.

			ICMc (DP)	
Network	I	L	α_{DP}	β
Full Last.fm	675 682	1 898 960	0.3	0.3
Last.fm USA	147 610	352 987	0.2	0.2

3.2 ICMc on Last.fm friendship network

Last.fm is an Internet site that learns the musical taste of its members on the basis of examples, and then constructs a personalized, radio-like music feed. The web site also has a richer array of services, including a possibility to announce friendships with other users. The friendships are initiated by a single party but are later mutual, forming a network with undirected links. Because friends tend to be similar, communities in the network would be relatively homogeneous by their musical taste and other characteristics. We use this similarity within communities to demostrate ICMc components.

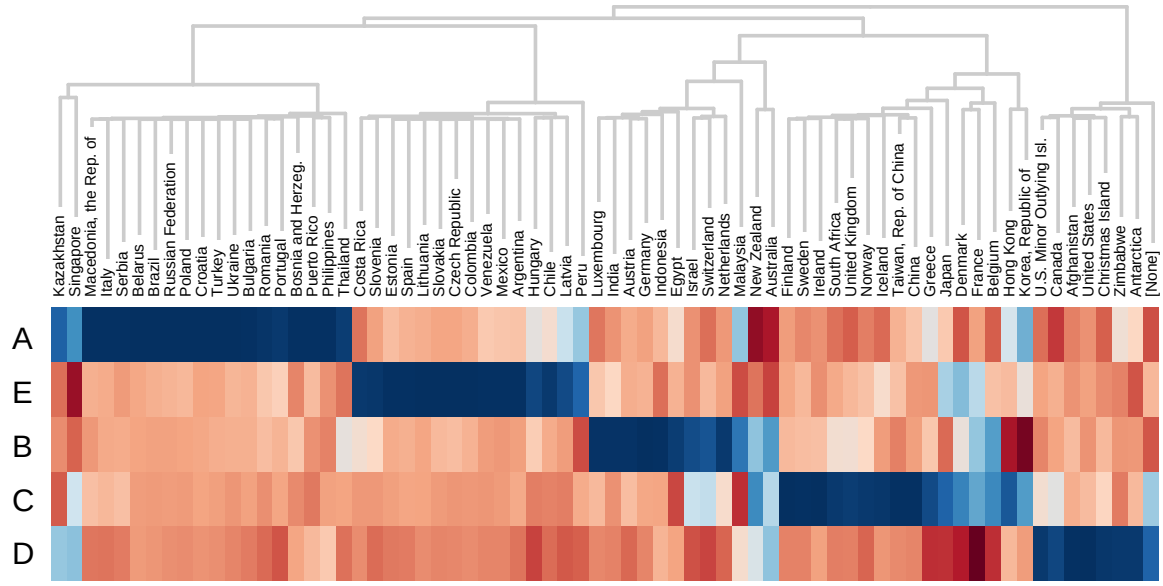

Figure 7: ICMc components (rows) of the Last.fm friendship network correlated with nationalities of the participants (columns). The full Last.fm network, about 675,000 nodes, was analyzed with the DP-version of ICMc. After a burn-in period of 19,000 iterations, 20 samples were taken at intervals of 50 iterations to get component memberships for the nodes. The running time was 16.4 hours when run in a single thread on an eight-core 2 GHz Intel Xeon. Dark blue and dark red denote the extremes of high and low co-occurrence counts, respectively. The columns are ordered and the tree produced by *heatmap* of the statistical environment *R*.

The global Last.fm network had about 675,000 nodes and 1,9 million links, while the subset of US members had about 147,000 nodes and 353,000 links (Table 2). In addition to the friendships, we also crawled the nationalities of the site members in the network, as well as the tags they had associated to the music they like. The most common tags represent musical genres or subgenres, allowing interpretation of the components found from the network.

We modeled the networks with the ICMc, with its Dirichlet process prior adjusted to favor few components. (With different hyperparameters, it would have been possible to obtain thousands of local communities, but the interpretation of such a solution here to get an idea about its quality would be difficult.) See Table 2 and Figs. 7 and 8 for details and results.

The component structure of the full Last.fm network is primarily about geography or nationalities (Fig. 7). This was unexpected at first sight, but in hindsight it is not at all surprising, for people tend to bond mostly within their country or city, and the friendships in Last.fm are likely to reflect the relationships of the real world. Even if they did not, nationality would affect bonding. We also correlated the global component structure to

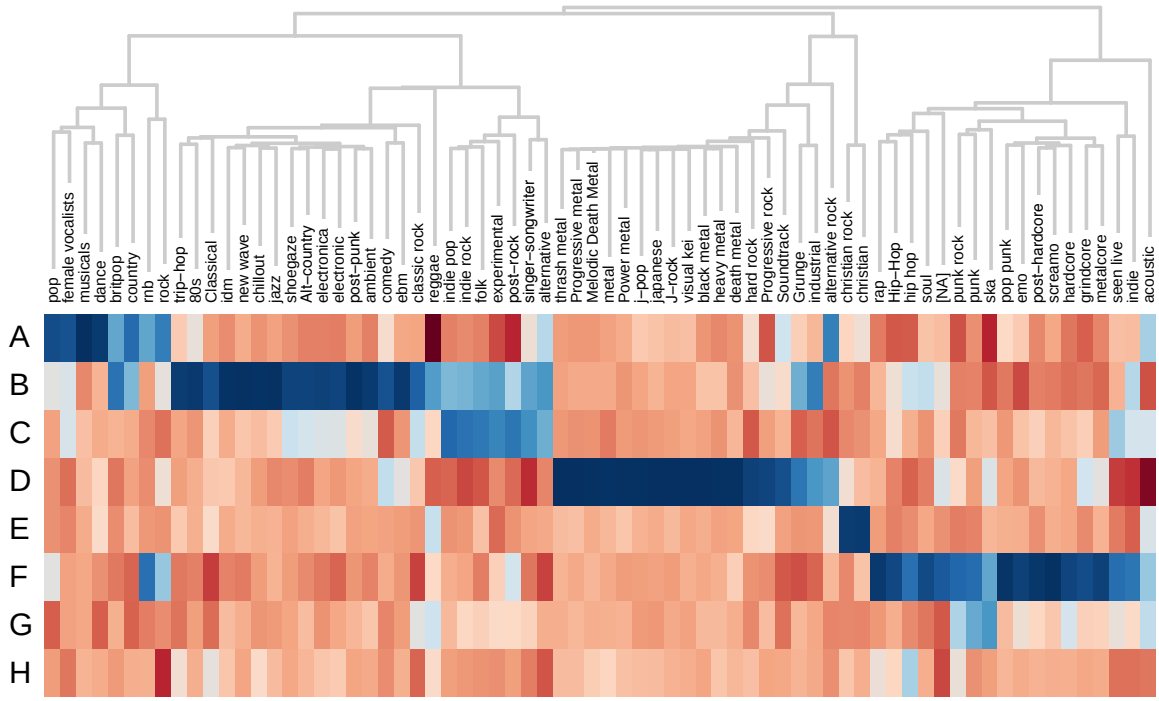

Figure 8: ICMc components (rows) of the Last.fm USA friendship network correlated with musical taste of the participants (columns). The model had the DP prior and was computed with a burn-in of 49,000 iterations, after which 20 samples were recorded at intervals of 50 iterations. The running time was 8.4 hours when run in a single thread on an eight-core 2 GHz Intel Xeon. Columns correspond to the most common tags given to the songs by the users themselves. Other details are as in Fig. 7.

musical taste, and while there are meaningful groups of genres (not shown, but see Fig. 8), it is hard to say which part of them arises due to the geographical division.

Although a more complex model would be needed to find both musical and geographical structure, the results show that ICMc is able to find homophilic structures from large networks. To get a better grasp of the musical homophily of the network, we also ran ICMc on a geographically more homogeneous subset of members who have announced to be from the US. This revealed a clear structure in terms of music preferences, as shown in Figure 8 and in Table 3. The model was able to separate light pop, more experimental music, "alternative," metal, Christian, and a punk–hip-hop continuum. In addition, there were two components that are harder to interpret.

4. Discussion

We have presented two generative models for networks, of which ICMc is novel, and demonstrated and tested them on data sets of various sizes. Performance differences between the

(a)

Cluster A	
juggalo	1.36
pop	1.34
musicals	1.32
Sludge	−1.89
black metal	−1.98

(b)

Cluster B	
shoegaze	1.35
Alt-country	1.24
post-punk	1.22
screamo	−1.79
pop punk	−3.16

(c)

Cluster C	
indie	0.46
post-rock	0.30
folk	0.22
visual kei	−1.86
j-pop	−2.08

(d)

Cluster D	
j-pop	1.69
visual kei	1.68
black metal	1.56
post-punk	−1.21
psychedelic	−1.41

(e)

Cluster E	
christian	1.53
podcast	1.01
trance	0.87
shoegaze	−1.54
Sludge	−1.68

(f)

Cluster F	
rnb	1.33
screamo	1.21
pop punk	1.15
Korean	−2.25
psytrance	−2.48

(g)

Cluster G	
Jam	1.35
ska	0.89
hardcore	0.47
visual kei	−1.43
j-pop	−2.28

(h)

Cluster H	
latin	1.13
chinese	1.05
psytrance	0.70
synthpop	−1.54
juggalo	−1.62

Table 3: The most likely and unlikely tags for each of the ICMc components in the Last.fm US network. The tables have been obtained by comparing the frequency of the tag to that expected in terms of its marginal probabilities. The table includes only tags for which the deviation from the expectation was reliable in terms of a binomial test (p=0.05). The numerical values are log-odds.

models were small; ICMc performed slightly better on networks with strong subgroup cohesion, while SSN-LDA had an edge in finding more disassortative network structures. If one is after communities in a social network, there are both theoretical and empirical reasons to prefer ICMc. The models do not have significant differences in implementation complexity or ease of use.

SSN-LDA can be seen as a further development of similar kinds of models earlier applied to text documents. On the other hand, it is also a generalization of the model by Newman and Leicht (2006), which interestingly shares notable similarity with earlier text document models (Hofmann, 2001). ICMc belongs to the same model family with LDA, but introduces a generative process that is more faithful to the idea of subgroup cohesion. An earlier formulation of subgroup cohesion is modularity (Newman and Girvan, 2004; Newman, 2006), for which ICMc or its likelihood could be seen as an alternative. It would be interesting to explore the relationship between these two, especially as our simulations show that in general modularity increases monotonically during a Gibbs run or saturates and only slightly decreases before convergence (Aukia, 2007).

Most of our tests were on networks with a known community structure, which allowed us to set the number of components in advance and use the Dirichlet prior. In preliminary tests we also tried Dirichlet process priors with these networks, but performance was naturally worse since they did not have the prior knowledge about the expected ("known") component number. Another reason for the worse performance of DPs probably is that the size distribution of the communities is artificially even, for two related reasons. First, the networks have survived the selection process of becoming de factor standards for model

testing. Second, in many cases the communities have been manually set up to make them maximally informative or otherwise handy. A small number of even-sized communities does not fit well with the Dirichlet process prior, which assumes either a small number of communities with rather unequal size, or a very large number with more equal size. It is likely that in real applications to social and biological networks the Dirichlet process performs relatively much better, because real-life communities tends to be of heterogenerous sizes.

The generative processes of simple models as discussed here are not meant to be realistic, at least not on higher hierarchical levels beyond the distributions generating the observed data. Instead, the ultimate criterion for generative processes should be empirical. Some abstract information about the networks can be coded to the generative processes, however. One obvious example is the assortative vs. disassortative nature of the network structure. It seems that getting this wrong is not catastrophical, but certainly using the right model improves performance. Another interesting detail are the Dirichlet priors. From their urn representations, it is obvious that they mimic the preferential attachment model of network generation (Albert and Barabasi, 2002) which produces relatively realistic degree distributions for social networks.

Even with the Dirichlet process prior one needs to choose the hyperparameters. Fortunately, the models seem to be quite robust in terms of the parameter β, and also in terms of α_{Dir} with the Dirichlet prior. It is possible to take β into the sampling process as an MCMC step, because the marginal likelihood for β is easy to compute. With the Dirichlet process prior, the parameter α_{DP} fundamentally affects the latent component diversity and therefore model complexity. For α one can use the proposed approximations of evidence, such as the harmonic mean estimator (Griffiths and Steyvers, 2004; Buntine and Jakulin, 2004)—it is known to be unstable but at least sometimes repairable (Raftery et al., 2007). Cross validation on the link level is still another possibility.

Although Gibbs sampling has a reputation of being slow compared to variational methods, a lot depends on how the slowness is measured. With topic models for texts, Gibbs is know to produce better results than variational LDA, at the cost of maybe 4–8 times the running time to convergence (Wray Buntine, p.c.). But according to Griffiths and Steyvers (2004), Fig. 1, collapsed Gibbs is actually *faster*, measured in floating point operations per second to attain a certain level of perplexity. The difference may partly be explained by implementational details, but one should also note that performance measurements should be relative to the goal: While in statistical inference convergence is essential, in predictive tasks the predictive performance counts, and often in practice a model is better if it gives better performance in a shorter running time, regardless of whether it has converged or not.

In fact, the whole notion of posterior convergence is problematic in models like LDA and ICMc with a high number of data, parameters and components. We do know that permutation modes exist and that the current Gibbs samplers fortunately find only one of them—if they found more, we would have a label switching problem. Even within a permutation mode there are probably many local modes of which the Gibbs sampler explores only part—this is suggested by the variation between the chains, and the NP-hardness of related formulations of the community finding problem (Brandes et al., 2006). If needed, different types of compromizes between running time and performance are available by applying better MCMC techniques, such as annealing, population methods, or split-merge

moves. Variational methods are available for the DP prior (Blei and Jordan, 2004) but they are likely to need help with mode finding.

ICMc and SSN-LDA can be considered as examples of a larger family of component models, giving generalizations. Links or higher-order co-occurrences of potentially several types are generated from latent components, together with other nominal data associated to nodes. Optimization of such models with collapsed Gibbs is relatively straightforward and easy to implement, as long as the priors are conjugate, non-parametric or not. An interesting extension of ICMc, evidently needed for the Last.fm network, would be to allow factorial (nominal) components, whose interactions describe the observed communities. In the Last.fm network, the obvious factors could be geography and musical taste. More generic formal extensibility of the model family, along the lines of relational models (e.g. Xu et al., 2007) should also be investigated.

Acknowledgments

We thank Last.fm for the data, and E. Airoldi for information on the block models, particularly on their scalability. This work was supported by Academy of Finland, grant number 119342. SK belongs to the Finnish CoE on Adaptive Informatics Research Centre of the Academy of Finland and Helsinki Institute for Information Technology, and was partially supported by EU Network of Excellence PASCAL.

Appendix A. The joint distribution and collapsed Gibbs sampler.

In ICMc the joint likelihood of observed links $\mathcal{L}$ and latent variables $\mathcal{Z}$, given mid-level model parameters θ and m, is

$$p(\mathcal{L}, \mathcal{Z}|m, \theta) = \prod_l \theta_{\mathcal{Z}_l} m_{\mathcal{Z}_l \mathcal{I}_l} m_{\mathcal{Z}_l \mathcal{J}_l} = \prod_z \theta_z^{n_z} \times \prod_{iz} m_{zi}^{k_{zi}} ,$$

where the notation $\mathcal{Z}_l$ refers to the index of the component generating link l, and $\mathcal{I}_l$ and $\mathcal{J}_l$ refer to link endpoint node indices. In the last expression we have link endpoints counts n_z over components, and k_{zi} over component–node co-occurrences. With symmetric Dirichlet priors $\mathrm{Dir}^I_{\mathrm{sym}}(\beta)$ for each m_z and $\mathrm{Dir}^K_{\mathrm{sym}}(\alpha_{Dir})$ for θ, this becomes

$$p(\mathcal{L}, \mathcal{Z}, m, \theta|\alpha_{Dir}, \beta) = Z^{-1}(\alpha_{Dir}, \beta) \prod_z \theta_z^{n_z + \alpha_{Dir} - 1} \prod_{iz} m_{zi}^{k_{zi} + \beta - 1} ,$$

with the normalizer Z arising from the Dirichlet priors. Following Griffiths and Steyvers (2004) on Rao-Blackwellisation of LDA, marginalize over θ and all m_z:

$$\begin{aligned} p(\mathcal{L}, \mathcal{Z}|\alpha_{Dir}, \beta) &= \iint p(\mathcal{L}, \mathcal{Z}, m, \theta|\alpha_{Dir}, \beta)\, d\theta\, dm \\ &= Z^{-1}(\alpha_{Dir}, \beta) \prod_z \frac{\prod_i \Gamma(k_{zi} + \beta)}{\Gamma(2n_z + M\beta)} \times \frac{\prod_z \Gamma(n_z + \alpha_{Dir})}{\Gamma(N + K\alpha_{Dir})} , \quad (12) \end{aligned}$$

where M is the number of nodes, K is the number of components, and the $2n_z$ comes from the number of component-wise links and the fact that each link has two endpoints. (For

evaluating the integral, look for a correspondence with the general Dirichlet distribution and its normalizing factor.)

Because links are generated independently, they can in principle be separated from $p(\mathcal{L}, \mathcal{Z}|\alpha, \beta)$ into link-wise factors. Separate one arbitrary link, say l_0, associated to the latent variable z_0 and to nodes i_0 and j_0 ($i_0 \neq j_0$), from the product, and denote by $(\mathcal{L}', \mathcal{Z}')$ the other links and their associated latent components, and by (k', n', N') the counts as they were if the link was nonexistent. For most indices, we will have $k' = k$ and $n' = n$, and always $N' = N - 1$, but for some indices $k' = k - 1$ and $n' = n - 1$. Because

$$\begin{aligned} \Gamma(x) &= (x-1)\Gamma(x-1) \\ \Gamma(x) &= (x-1)(x-2)\Gamma(x-2) \ , \end{aligned}$$

all this translates into

$$\begin{aligned} p(\mathcal{L}', \mathcal{Z}', l_0, z_0|\alpha_{Dir}, \beta) = Z^{-1}(\alpha_{Dir}, \beta) \prod_z \frac{\prod_i \Gamma(k'_{zi} + \beta)}{\Gamma(2n'_z + M\beta)} \times \frac{\prod_z \Gamma(n'_z + \alpha_{Dir})}{\Gamma(N' + K\alpha_{Dir})} \times u_0 \\ = p(\mathcal{L}', \mathcal{Z}'|\alpha_{Dir}, \beta) \times u_z \ , \end{aligned}$$

where

$$u_z \equiv p(l_0, z_0|\mathcal{L}', \mathcal{Z}', \alpha_{Dir}, \beta) = \frac{(k'_{z_0 i_0} + \beta)(k'_{z_0 j_0} + \beta)}{(2n'_{z_0} + 1 + M\beta)(2n'_{z_0} + M\beta)} \times \frac{n'_{z_0} + \alpha_{Dir}}{N' + K\alpha_{Dir}} \ . \qquad (13)$$

One can use the result to sample a new component z for the left-out link, with the probabilities $p(z|l_0, \mathcal{L}', \mathcal{Z}', \alpha_{Dir}, \beta) = u_z/u.$, the denominator using the dot notation for the sum. A Gibbs iteration follows by leaving one link out at a time, and sampling a new latent component for it as above.

Dirichlet process prior for components. The ICMc model can be derived for a Dirichlet Process component prior in several ways. Informally, after seeing the link removal decomposition with u_z, one notes the structure of $p(\mathcal{L}, \mathcal{Z}|\alpha_{Dir}, \beta)$ as nested Polya urns (Johnson, 1977). One can then substitute the component urn, the last factor in (13), with the Blackwell-MacQueen urn (Blackwell and MacQueen, 1973; Tavare and Ewens, 1997) parameterized by α_{DP}:

$$p(l_0, z_0|\mathcal{L}', \mathcal{Z}', \alpha_{DP}, \beta) = \frac{(k'_{z_0 i_0} + \beta)(k'_{z_0 j_0} + \beta)}{(2n'_{z_0} + 1 + M\beta)(2n'_{z_0} + M\beta)} \frac{\mathrm{C}(n_{z_0}, \alpha_{DP})}{N + \alpha_{DP}} \qquad (14)$$

with $\mathrm{C}(n, \alpha) = n$ if $n \neq 0$ and $\mathrm{C}(0, \alpha) = \alpha$.

Another way to end up with the same result is to substitute $\alpha_{Dir} = \alpha_{DP}/K$ to (12) or (13), then collect all empty components into one bin, and take the limit $K \to \infty$ (Neal, 2000).

More formally, one can first write the joint distribution of the ICMc model with an unspecified component prior $p(\mathcal{Z}|\alpha)$,

$$p(\mathcal{L}, \mathcal{Z}, m|\alpha, \beta) = p(\mathcal{L}, m|\mathcal{Z}, \beta)\, p(\mathcal{Z}|\alpha) = Z^{-1}(\beta) \prod_{iz} m_{zi}^{k_{zi}+\beta-1} \times p(\mathcal{Z}|\alpha) \ ,$$

integrate m out, and then substitute the Dirichlet process prior (e.g., Dahl 2003), obtainable from the Blackwell-Queen urn model by induction, to end up with

$$p(\mathcal{L}, \mathcal{Z}|\alpha_{DP}, \beta) = Z^{-1}(\beta) \prod_z \frac{\prod_i \Gamma(k_{zi} + \beta)}{\Gamma(2n_z + M\beta)} \times \frac{\alpha_{Dir}^K \Gamma(\alpha_{Dir})}{\Gamma(\alpha_{Dir} + N)} \prod_z \Gamma(n_z) \; . \tag{15}$$

The sampling rule (14) can then be obtained by computing the probability of one (removed) link given all others, just as in the case of a finite Dirichlet prior.

Collapsed Gibbs sampling for the SSN-LDA model. The collapsed sampler is identical to that in Griffiths and Steyvers (2004), also presented by Zhang et al. (2007). The collapsed sampling formula for SSN-LDA with the DP prior is obtained analogously to ICMc, by modifying the factor corresponding to the latent-component urn.

References

E. M. Airodi, D. M. Blei, S. E. Fienberg, and E. P. Xing. Mixed membership stochastic blockmodels. *Journal of Machine Learning Research*, 2008. In press.

Reka Albert and Albert-Laszlo Barabasi. Statistical mechanics of complex networks. *Reviews of Modern Physics*, 74(1):47–97, 2002.

Janne Aukia. Bayesian clustering of huge friendship networks. Master's thesis, Laboratory of Computer and Information Science, Helsinki University of Technology, Espoo, Finland, 2007.

D. Blackwell and J. B. MacQueen. Ferguson distributions via Polya urn schemes. *Annals of Statistics*, 1:353–355, 1973.

D. Blei, A. Ng, and M. Jordan. Latent Dirichlet allocation. *Journal of Machine Learning Research*, 3:993–1022, 2003.

D. M. Blei and M. I. Jordan. Variational methods for the Dirichlet process. In *Proceedings of the 21st International Conference on Machine Learning (ICML)*, 2004.

James L. Blue, Isabel Beichl, and Francis Sullivan. Faster Monte Carlo simulations. *Physical Review E*, 51(2):867–868, 1995.

Ulrik Brandes, Daniel Delling, Marco Gaertler, Robert Grke, Martin Hoefler, Zoran Nikoloski, and Dorothea Wagner. Maximizing modularity is hard. *ArXiv e-prints*, 2006. arXiv:physics/0608255v2.

W. Buntine and A. Jakulin. Applying discrete PCA in data analysis. In M. Chickering and J. Halpern, editors, *Proc. UAI'04, 20th Conference on Uncertainty in Artificial Intelligence*, pages 59–66. AUAI Press, 2004.

Wray L. Buntine. Variational extensions to EM and multinomial PCA. In Tapio Elomaa, Heikki Mannila, and Hannu Toivonen, editors, *Proceedings of the 13th European Conference on Machine Learning (ECML'02)*, volume 2430 of *Lecture Notes in Computer Science*, pages 23–34, London, UK, 2002. Springer.

J. J. Daudin, F. Picard, and S. Robin. A mixture model for random graphs: A variational approach. Technical Report 4, Statistics for systems biology group, INRA, Jouy-en-Josas, France, 2007.

Santo Fortunato and Marc Barthelemy. Resolution limit in community detection. *Proceedings of the National Academy of Sciences USA*, 104(1):36–41, 2007.

Santo Fortunato and Claudio Castellano. Community structure in graphs. *ArXiv e-prints*, 2007. arXiv:0712.2716.

S. Geman and D. Geman. Stochastic relaxation, Gibbs distributions, and the Bayesian restoration of images. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 6(6):721–741, 1984.

M. Girvan and M. E. J. Newman. Community structure in social and biological networks. *Proceedings of the National Academy of Sciences USA*, 99(12):7821–7826, 2002.

Thomas L. Griffiths and Mark Steyvers. Finding scientific topics. *Proceedings of the National Academy of Sciences USA*, 101 Suppl 1:5228–5235, 2004.

Mark S. Handcock, Adrian E. Raftery, and Jeremy M. Tantrum. Model-based clustering for social networks. *Journal Of The Royal Statistical Society Series A*, 170(2):301–354, 2007.

Jake M. Hofman and Chris H. Wiggins. A Bayesian approach to network modularity. *ArXiv e-prints*, 2007. arXiv:0709.3512v2.

Thomas Hofmann. Unsupervised learning by probabilistic latent semantic analysis. *Machine Learning*, 42:177–196, 2001.

N. L. Johnson. *Urn models and their applications*. John Wiley and Sons, 1977.

Jussi M. Kumpula, Jari Saramaki, Kimmo Kaski, and Janos Kertesz. Limited resolution in complex network community detection with Potts model approach. *The European Physical Journal B*, 56(1):41–45, 2007.

Paul Lazarsfeld and Robert K. Merton. Friendship as a social process: A substantive and methodological analysis. In *Freedom and Control in Modern Society*, pages 18–66. Van Nostrand, New York, USA, 1954.

Alaina Michaelson and Noshir S. Contractor. Structural position and perceived similarity. *Social Psychology Quarterly*, 55(3):300–310, 1992.

Radford M. Neal. Markov chain sampling methods for Dirichlet process mixture models. *Journal of Computational and Graphical Statistics*, 9(2):249–265, 2000.

M. E. J. Newman. The structure and function of complex networks. *SIAM Review*, 45(2): 167–256, 2003.

M. E. J. Newman. Modularity and community structure in networks. *Proceedings of the National Academy of Sciences USA*, 103:8577–8582, 2006.

M. E. J. Newman and M. Girvan. Finding and evaluating community structure in networks. *Physical Review E*, 69(2):026113, 2004.

M. E. J. Newman and E. A. Leicht. Mixture models and exploratory data analysis in networks. *ArXiv e-prints*, 2006. arXiv:physics/0611158.

A. E. Raftery, M. A. Newton, J. M. Satagopan, and P. Krivitsky. Estimating the integrated likelihood via posterior simulation using the harmonic mean identity (with discussion). In J. M. Bernardo et al., editor, *Bayesian Statistics*, volume 8, pages 1–45. Oxford University Press, 2007.

Prithviraj Sen and Lise Getoor. Link-based classification. Technical Report Report CS-TR-4858, University of Maryland, College Park, USA, 2007.

Janne Sinkkonen, Janne Aukia, and Samuel Kaski. Inferring vertex properties from topology in large networks. In *Working Notes of the 5th International Workshop on Mining and Learning with Graphs (MLG'07)*, Florence, Italy, 2007. Universita degli Studi di Firenze. Extended Abstract.

S. Tavare and W. J. Ewens. The Ewens sampling formula. In *Multivariate discrete distributions*. John Wiley & Sons, New York, USA, 1997.

S. Wasserman and K. Faust. *Social Network Analysis: Methods and Applications*. Cambridge University Press, Cambridge, UK, 1994.

Mark Allen Weiss. *Data Structures and Algorithm Analysis in Java*. Addison-Wesley, Reading, USA, 1998.

C. K. Wong and M. C. Easton. An efficient method for weighted sampling without replacement. *SIAM Journal on Computing*, 9(1):111–113, 1980.

Zhao Xu, Volker Tresp, Shipeng Yu, Kai Yu, and Hans-Peter Kriegel. Fast inference in infinite hidden relational models. In *Working Notes of the 5th International Workshop on Mining and Learning with Graphs (MLG'07)*, Florence, Italy, 2007. Universita degli Studi di Firenze. Extended Abstract.

Wayne W. Zachary. An information flow model for conflict and fission in small groups. *Journal of Anthropological Research*, 33(4):452–473, 1977.

Haizheng Zhang, Baojun Qiu, C. Lee Giles, Henry C. Foley, and John Yen. An LDA-based community structure discovery approach for large-scale social networks. In *Intelligence and Security Informatics (ISI) 2007*, pages 200–207. IEEE, 2007.